Radiative transfer for cloud-scale models: accuracy and efficiency

Robert Pincus and Bjorn Stevens

Terminology

By “cloud-scale model” I mean models that resolve the circulations in individual clouds

Roughly:

“large-eddy simulations” grid scale O(10 m)

“cloud-resolving models” grid scale O(1 km)

Circulations are driven by internal heating/cooling and surface fluxes

Radiation for cloud scale models: a perfect world

The "proper" radiation calculation is broadband 3D radiative transfer at each time step, but
a) this is horribly expensive, and
b) heating rate differences from 1D are small

Next easiest is independent broadband 1D calculations in each column ("ICA") at each time step

$$F(x,y,t) = \sum_{b}^{B} w_b \sum_{g}^{G(b)} w_{g(b)} F_{b,g}(x,y,t)$$

But even this isn't practical: our naïve implementation increased solution time by a factor of 50

Radiation for cloud-scale models: the real world

But radiation is often a small forcing at the cloud scale
(not always, though - think stratocumulus!)

So in cloud-scale models, radiation may be

ignored

idealized

parameterized simply (i.e. analytic fits)

or, for the most flexibility (think MMF)

computed every *N* times steps

(GCMs do this too)

Why infrequent radiation calculations are a bad idea

The choice of *N* is arbitrary: no objective convergence tests

No way to know when *N* is too big
(and some systems are known to be unstable)

Sampling errors are correlated with the flow
(increase with local velocity scale)

So we tried another approach (stop me if you've heard this)

Why are heating rate calculations so expensive? It's the broadband integration - the double sum over bands and *g*-points

So we turned the roadblock into a springboard:

So we tried another approach (stop me if you've heard this)

Why are heating rate calculations so expensive? It's the broadband integration - the double sum over bands and *g*-points

So we turned the roadblock into a springboard:

$$F(x,y,t) \approx F_{MC}(x,y,t) = w(b')F_{b',g'}(x,y,t)$$
$$\text{where}$$
$$p(b') = 1/B \text{ and } p(g') = 1/w_{g'(b')}$$

Formally, this is a Monte Carlo sample of the full calculation, so we're calling this "Monte Carlo spectral integration"

A single estimate is noisy but many estimates converge to the right answer.

We built this. It works like a champ

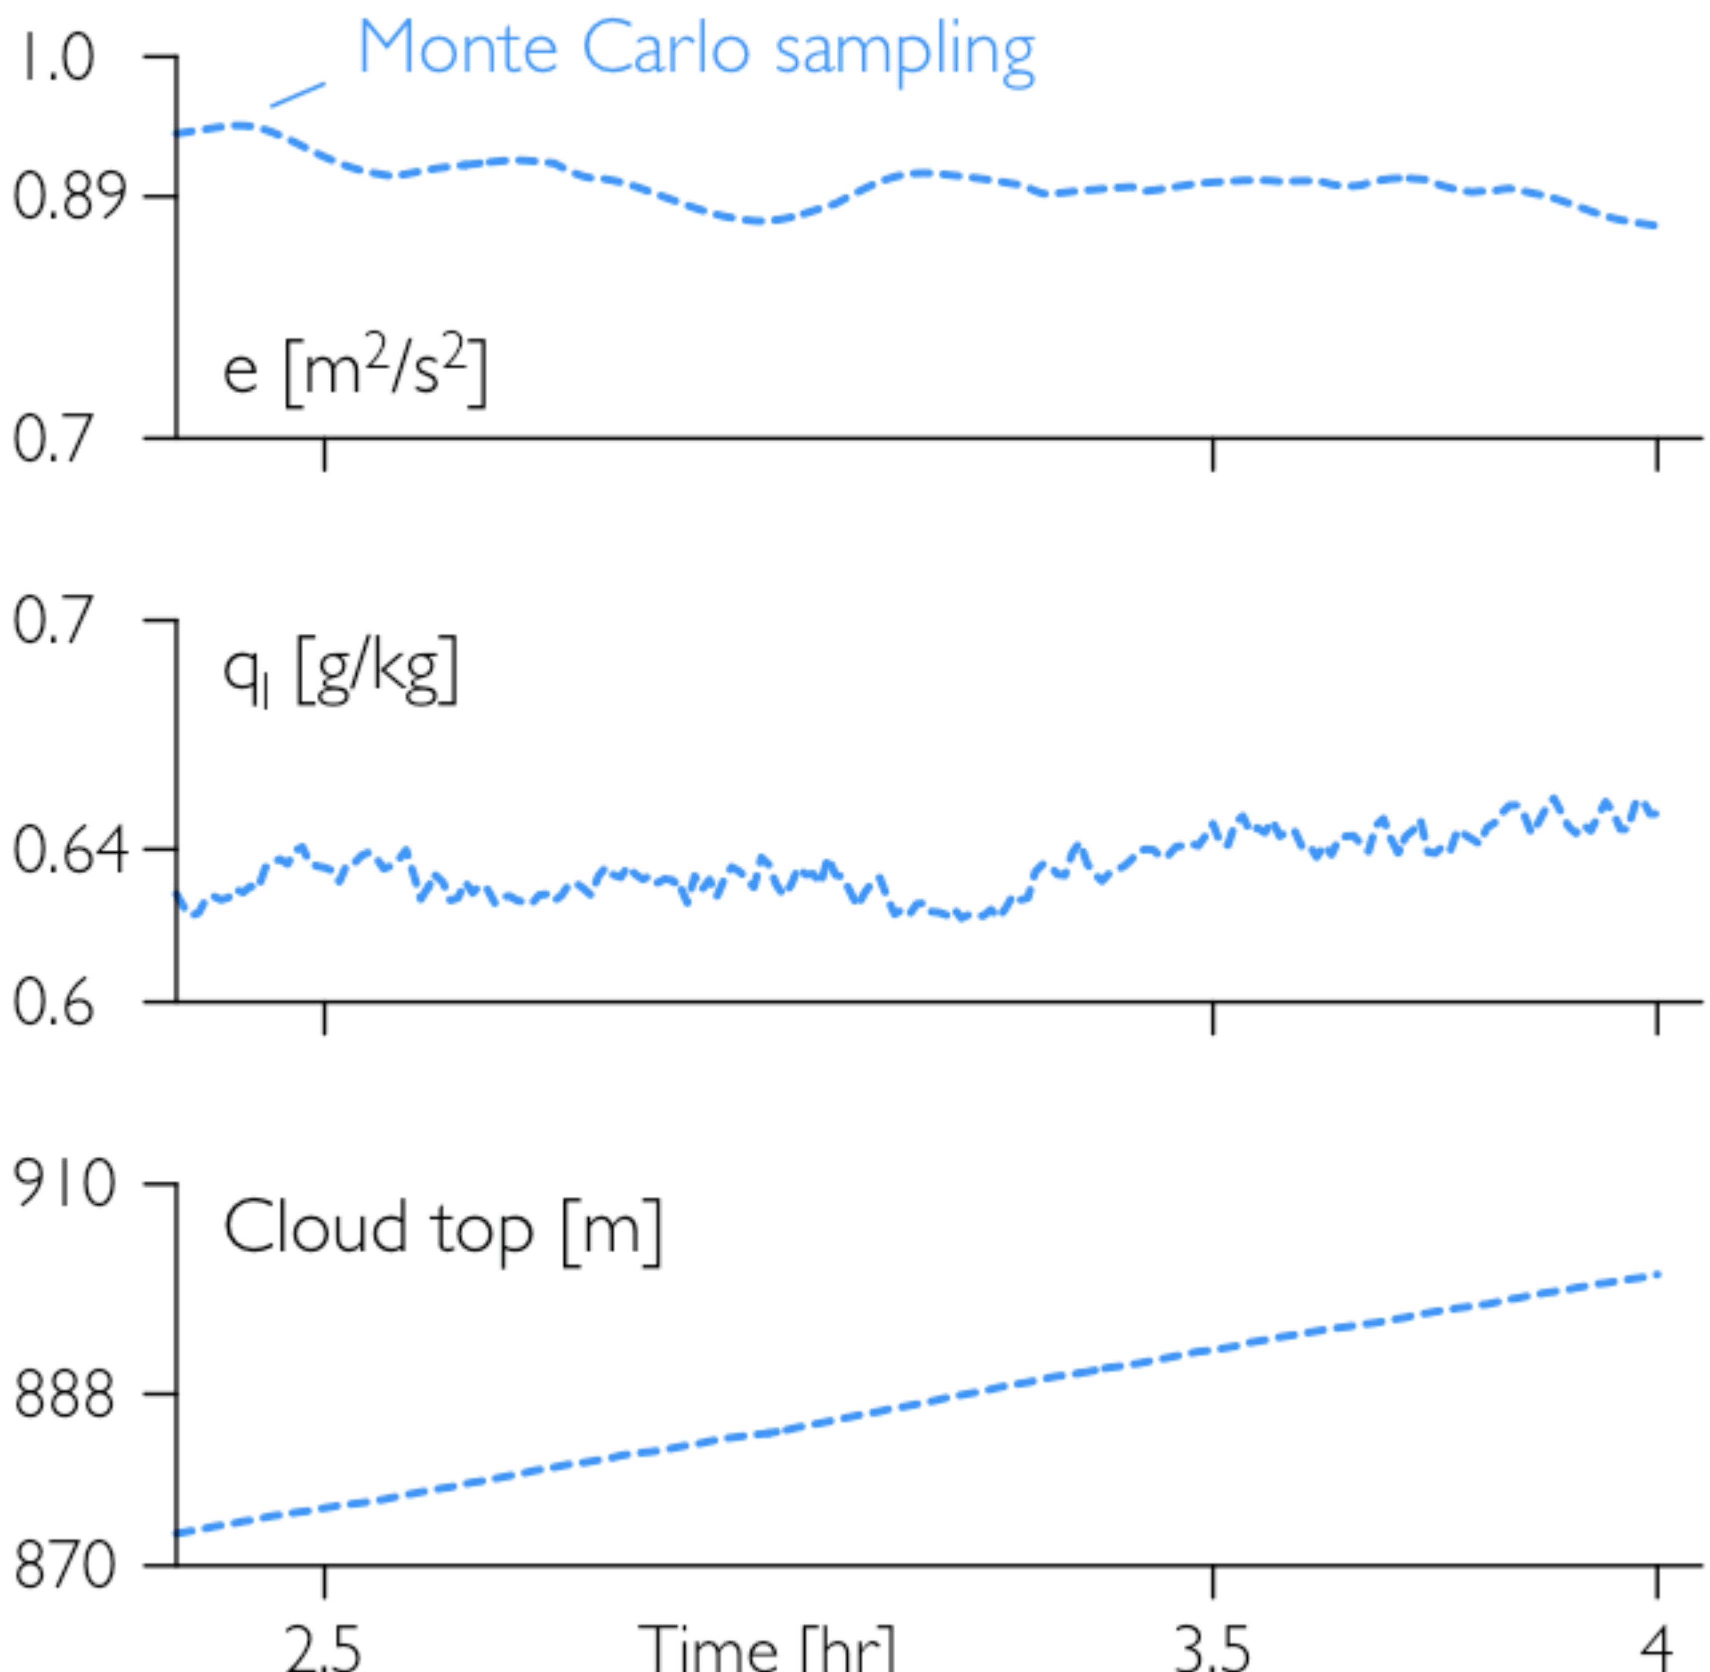

Pincus and Stevens, 2008

We built this. It works like a champ

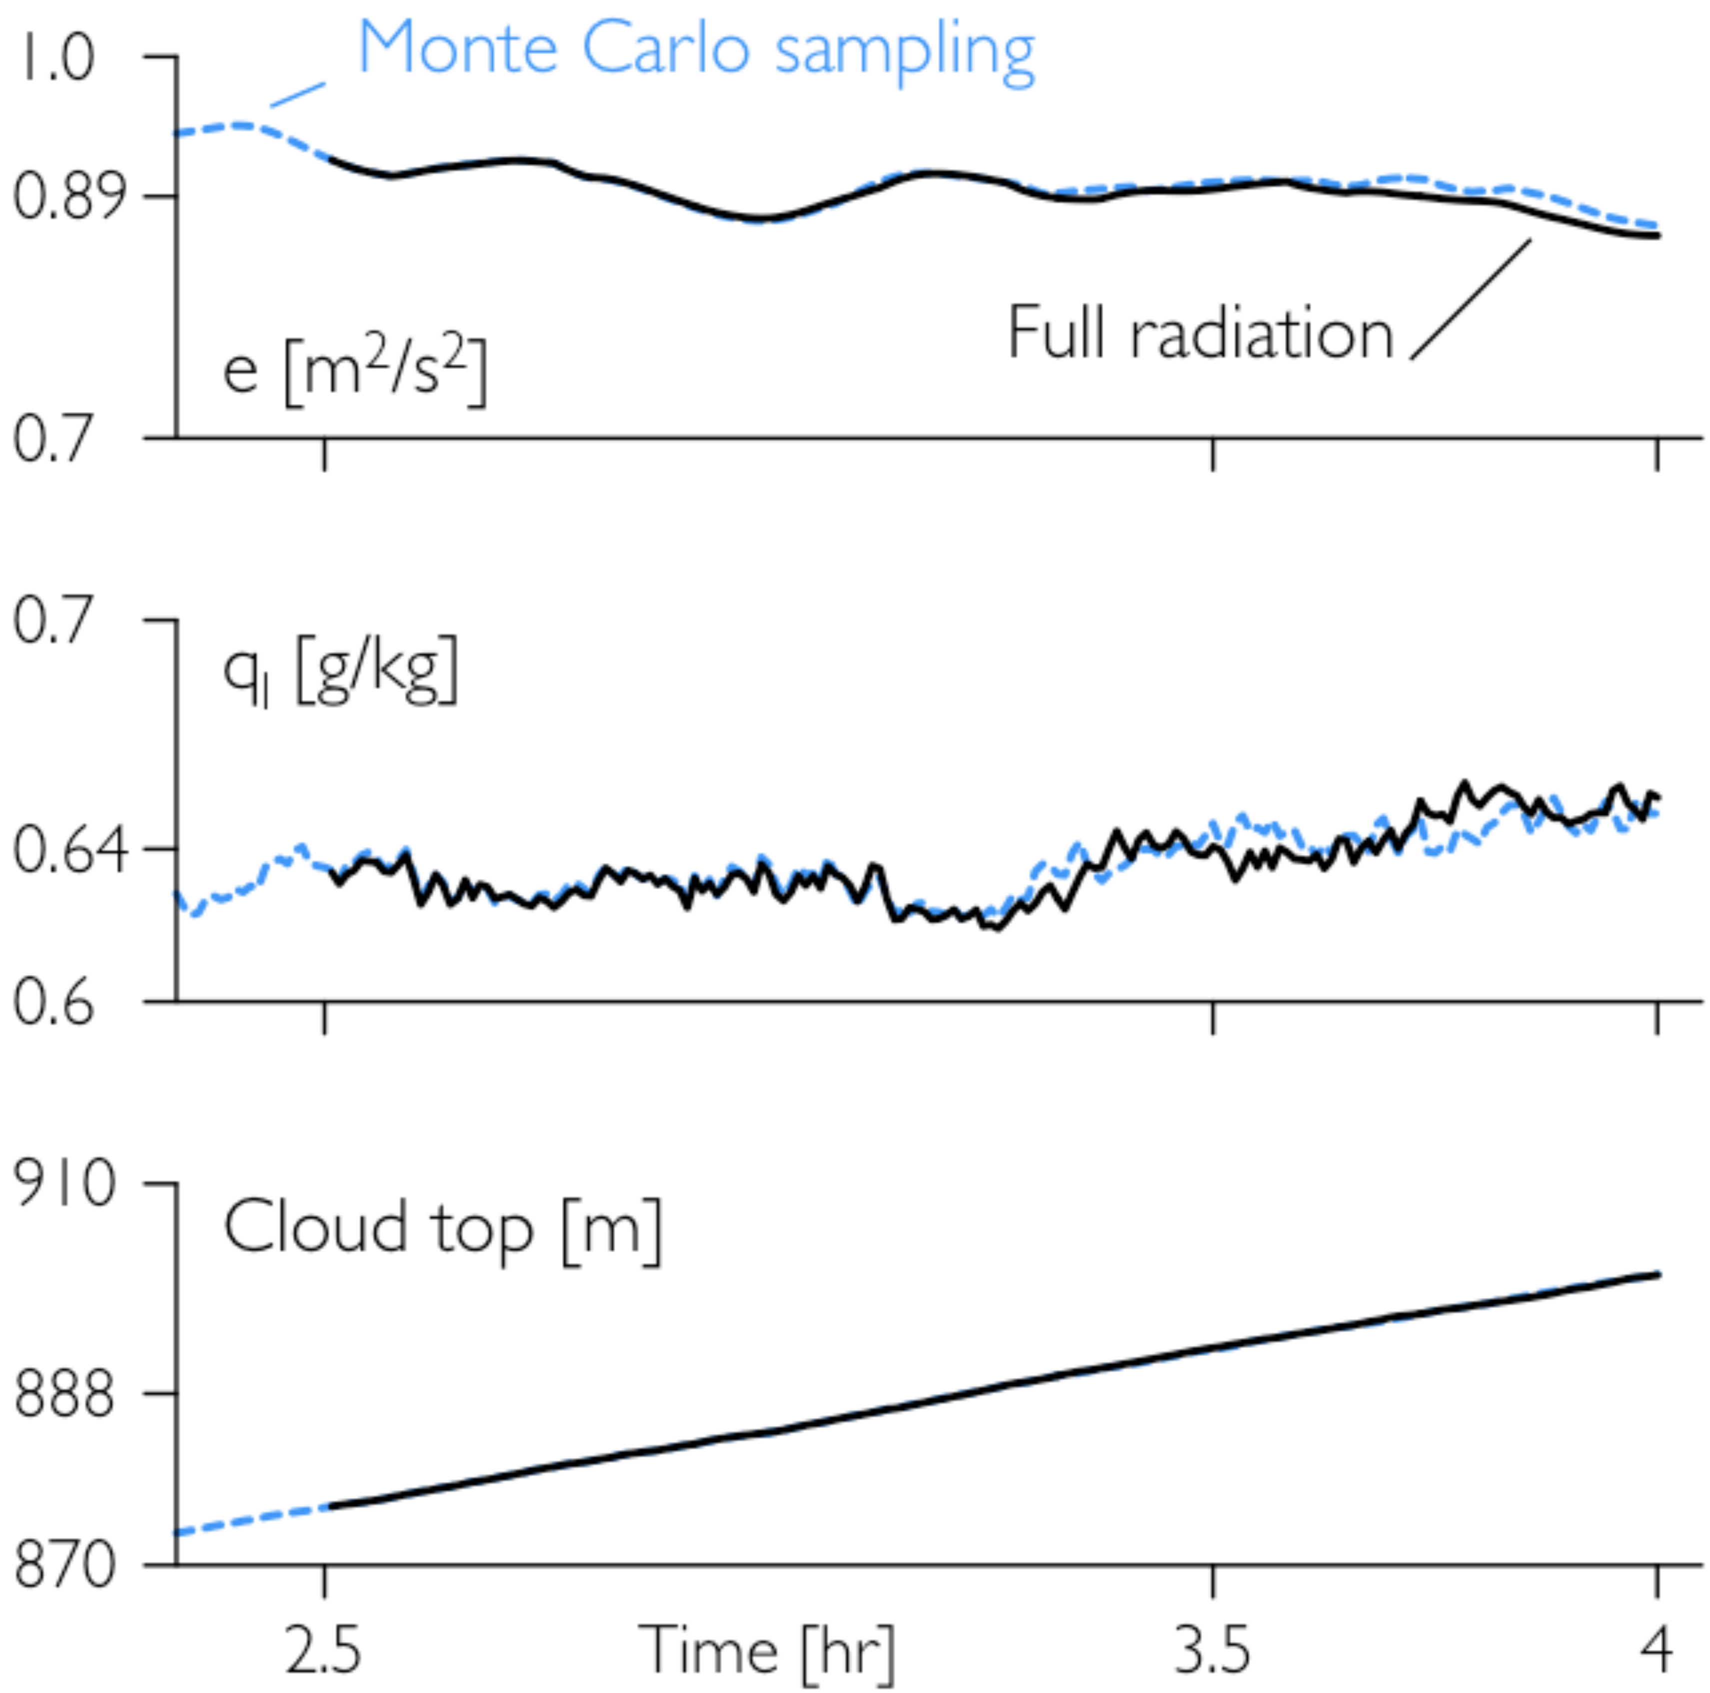

Pincus and Stevens, 2008

A scaling analysis for large-eddy simulation (i)

We'll compare the energy we expect in an eddy of a given size with the energy introduced by the Monte Carlo noise

Consider a well-mixed boundary layer with simple physics:

Radiative cooling at cloud top causes a buoyancy flux B_h

which drives eddies about as big as the boundary layer depth.

These drive smaller eddies according to the Komolgorov cascade

$$B_h \propto \Delta F$$
$$\overline{e}_h \propto (B_h h)^{2/3}$$
$$\overline{e}_l \propto \overline{e}_h \left(\frac{l}{h}\right)^{2/3}$$

A scaling analysis for large-eddy simulation (ii)

Now imagine some approximation to the driving fluxes

$$B_l = \overline{B_l} + B_l'$$

These perturbations systematically affect the flow if
1) they persist for an eddy turnover time, and
2) the perturbation changes the eddy energy significantly

A scaling analysis for large-eddy simulation (ii)

Now imagine some approximation to the driving fluxes

$$B_l = \overline{B_l} + B_l'$$

These perturbations systematically affect the flow if
1) they persist for an eddy turnover time, and
2) the perturbation changes the eddy energy significantly

Monte Carlo estimates of flux produce perturbations that scale as

$$B_l' = \overline{\sigma_B} / \sqrt{n_l}$$

n_l is determined by the scale l of the eddy, through the spatial scale and the CFL criteria

$$n_l = n_{l,xy} n_{l,t} \approx \left(\frac{l}{\delta x}\right)^2 \frac{\tau_l}{\delta t} \approx \left(\frac{l}{\delta x}\right)^3 \left(\frac{h}{l}\right)^{1/3}$$

A scaling analysis for large-eddy simulation (iii)

For resolved eddies $n_l >> 1$

$$e'_l = (\mathcal{B}'_l l)^{2/3} = \left(\frac{\sigma_{\mathcal{B}}}{\sqrt{n_l}} l\right)^{2/3} = (\sigma_{\mathcal{B}} l)^{2/3} \, \frac{\delta x}{l} \, \left(\frac{l}{h}\right)^{1/9}$$

So the ratio of the specific energy to the expected value is small:

$$\frac{e'_l}{\overline{e}_l} \propto \left(\frac{\sigma_B}{\overline{B}_l}\right)^{2/3} \frac{\delta x}{l} \left(\frac{l}{h}\right)^{1/9}$$

In (other) words

The technique introduces lots of noise, but that noise is

largest at the smallest time/space scales
(where it diffuses away quickly)
and
small relative to the energy from other source at resolved scales

Two implications

The practical:
We have a way to compute interactive radiation in cloud-scale models

The theoretical:
We have a way to understand how approximations for radiative transfer (think 3D vs 1D radiative transfer) affect simulations by cloud-scale models

I'll bet anyone a dollar we can get away with 1D radiative transfer in the shortwave